Soft Bound

ARTIFICIAL INTELLIGENCE AND EXPERT SYSTEMS

U-SAI
BACKUP

U-SAI

ARTIFICIAL INTELLIGENCE AND EXPERT SYSTEMS

CONTENTS

Page

Introduction 1

I ????

II EXECUTIVE SUMMARY 3

A. ~~Certainty and Uncertainty~~ Attitudes Toward New Technologies 4

B. Data/Information/Knowledge and 6

C. Don't Underrate Problems of AI

D. Don't Undervalue Concepts 10

E. A Positive Course of Action 12

III WHAT IS ARTIFICIAL INTELLIGENCE? 15

A. A Brief History 15

B. Artificial Intelligence Defined 18

C. AI OR/CC and 24

D. OR

IV KNOWLEDGE-BASE/EXPERT SYSTEMS 27

A. Status and Terminology 27

B. How Expert Systems Are Different 34

C. Hardware/Software Tools and 34

1. Problems of Implementation 38

2. LISP Machines and LISP

3. Expert System Software 42

D. Applications Areas 47

V INPUT'S ANALYSIS OF WHERE THINGS REALLY STAND

A. ~~Where Things Really Stand~~

B. 1. A Simple Case Study

C. 2. Expert System Fragility

D. 3. Ignorance and Expert Systems

E. 4. What Is Going to Happen?

F. 5. What You Should Not Do

G. 6. What You Should Do

D. Contraversy Decision and Game Theory

V INPUT'S ANALYSIS 51

A. Where Things Really Stand 51

1. A Simple Case Study 51

2. Expert Systems Fragility 54

3. Ignorance and Expert Systems 57

4. Decision Theory and Game Theory Contraversy 61

B. Recommendations 62

1. What Is Going to Happen? 62

2. What You Should Not Do 64

3. " " " Do 65

ARTIFICIAL INTELLIGENCE AND EXPERT SYSTEMS

EXHIBITS

			Page
II	-1	~~Certainty and Uncertainty~~ Attitudes toward New Technologies	5
	-2	Data, Information, and Knowledge	7
	-3	Don't Underrate Problems	9
	-4	Don't Undervalue AI Concepts	11
	-5	A Positive Course of Action	13
IV	-1	An Expert System Supported by AI	31
	-2	Distinguishing Characteristics of Expert Systems	35
	-3	The Components of an Expert System	39
	-4	Expert System Hardware	43
	-5	Intelligent Problem Solving	44
	-6	Expert System Software Tools ~~Generators~~	48
	-7	Categories of Expert Systems and Representative Impacts	49

U-SAI-I

1 ~~INTRODUCTION~~

o The fascination with artificial intelligence (AI) predates the electronic computer. In the 1800s a mechanical chess-playing robot was developed and enjoyed some brief success until it was discovered that a midget was concealed inside the mechanism. ~~While~~ Although advances in electronic technology have permitted increasing sophistication in hiding the human in the mechanism, there is still a substantial carnival element of conning the rules associated with the application of computer technology--especially in the adaptation of terminology. As one software manager was quoted in a recent issue of the Applied Artificial Intelligence Reporter: "Before long, anything that contains a slick algorithm will be declared to possess AI."

o INPUT's poll to determine what would be covered in this report was unanimous on two points. Users want ~~in stating that~~ definitions, and an understanding ~~were needed~~ of what constitutes AI, ~~and~~ specifically how expert or knowledge-based systems ~~were~~ are "different" ~~from what was already being done in some advanced planning and operational systems which already existed.~~ from conventional data processing systems.

o In addition, several very pragmatic questions were raised such as: "How much ~~is it (expert systems)~~ be ~~cost~~ justified?" Both are very good questions because ~~especially since nobody really understand what all the excitement is about except that Future and Business Week are writing articles about expert systems and management interest cannot be far behind.~~ most of the current excitement is being generated ~~by the press~~ not by users but by the media. "The trend toward expert systems going to cost ~~more~~?" and "How can ~~they~~ the cost of ~~a~~ expert systems

Stet

- The purpose of this study is to answer ~~such general~~ these questions ~~to the best of our ability~~ with special emphasis upon anticipated problems. Areas to be covered included:

  - AI and expert systems ~~-~~ definitions of
  - AI and operations research (OR) ~~-~~ affinities.
  - Hardware, software, and human ~~-~~ considerations ~~and~~ -- including costs.
  - Quality assurance (QA) ~~-~~ implications.
  - Applications and domains ~~-~~ ~~explained.~~ Explanation of
  - Information systems ~~-- does and don'ts. What to do, and what not to do~~ Integration ~~of data processing~~ activities.

- The methodology employed in developing this report was to review ~~as much~~ AI information ~~as time would permit~~, talk ~~to a few~~ with knowledgable people, and to reach intuitive conclusions. In other words, this product contains all natural intelligence without any artificial ~~flowing~~ flavoring, preservatives, or supplements. Taken with reasonable exercise of your own natural intelligence, it should ~~be~~ beneficial ~~to~~ both your personal and your corporate health and welfare.

2 main bullets

- Obviously, with such a wide range of topics it will be impossible to explore specific areas in any detail. Indeed, to do so would be misleading, because ~~there~~ substantial controversy remains concerning many major issues.
- You will find that this report has more of an academic orientation than most INPUT reports. This is true not only because AI is in the process of emerging from an academic environment, but also because there are many fundamental technical and ~~philosophical~~ philosophical questions about AI that ~~which have~~ remain unanswered.

II

II EXECUTIVE SUMMARY

- This executive summary is designed in presentation format in order to:

  - Help the reader review the key research findings.

  - Provide a general presentation and script to facilitate group communications.

- The key points of the entire report are summarized in Exhibits II-1 through II-5. On the left-hand page facing each exhibit is a script explaining the exhibits' contents.

- The entire field of Artificial Intelligence is shrouded in controversy, and research and development is bounded by the loading edge of many disciplines, including computer ~~service~~ science. In many ways this report is like the first version of an expert system, and the executive summary represents only the surface of a complex subject ~~—~~ the reader is ~~urged~~ expected to understand only the bounds of ~~its~~ the AI domain.

A. ~~CERTAINTY AND UNCERTAINTY~~ ATTITUDES TOWARD NEW TECHNOLOGIES

- When Charles Babbage conceived the Analytical Engine, ~~Lady~~ the Countess of Lovelace (who interpreted Babbage's work) observed that there was a natural tendency to first overrate and then undervalue all such developments. This is perhaps the most profound observation ever made concerning computer technology.

- Consider all the solutions and final answers two systems development ~~such as~~ ③ They include COBOL, Data Base Management Systems, Structured Programming, Fourth Generation Languages, etc. all have gone, ~~and~~ or are going, through some phase of Lady Lovelace's cycle.

- It is a certainty that artificial intelligence and expert systems will be subject to the same cycle, but uncertainty arises in determining exactly where we are in the cycle.

- Expert systems, unlike traditional data processing systems, deal with uncertainty--with those "maybe--I wonder" problems ~~which~~ that have no specific solution rather than the straightforward simplicity of accounting type problems ~~which~~ lead themselves to "yes or no" solutions.

← use ellipses ...

- Confronted with uncertainty, it is certain that both humans and expert systems will reach varrying solutions to the same problem. Establishing probabilities associated with the quality of such solutions is ~~certainly~~ a difficult problem.

B. DATA/INFORMATION/KNOWLEDGE

- Data are facts and a great deal of progress has been made in storing data on electronic media for processing and access by computers.

- Information is the result of analysis and may be based on ~~both~~ facts and/or opinions. It is normally stored on paper--~~in either~~ file cabinets or libraries. Office automation is concerned primarily with the movement of information from paper to electronic media where it can be more readily processed, formatted, and accessed. Despite some progress, most information continues to be stored on paper.

- Knowledge is the result of human analysis of both data and information. The process of human knowledge acquisition (learning) and storage (memory) are not very well understood. Confronted with a vast array of data, information, and uncertainty, humans have the ability to solve problems (and make decisions) with varying degrees of success.

- At present, computer information systems facilitate making data and information more readily available to human beings, but human knowledge is required for practically all complex problem solution and decision making. Expert systems are attempting to put human knowledge in electronic form (knowledge bases), and the process of doing this is referred to as knowledge engineering. (This is represented by the broken-line circle on the Data/Information/Knowledge exhibit.)

- Conceptually, data, information and knowledge would converge into intelligent electronic systems, but progress has been excrutiatingly slow. Despite considerable publicity, there are significant barriers to the development of intelligent machines for practical problem solving.

(This would mean that human knowledge would attach and store in electronic form along with data and information)

C. DON'T UNDERRATE PROBLEMS

- The primary reason ~~Lady~~ the Lovelace's cycle exists is because the computer industry has traditionally underrated implementation problems. As artificial intelligence and expert systems emerge from the academic environment, there are indications that implementation problems will be especially severe.

- There is fundamental ignorance concerning how humans learn, how knowledge is stored, and ~~the role~~ how intuition ~~plays in~~ affects complex problem ~~solution~~ solving. These difficult problems are being disregarded by many AI practitioners.

- The AI community has isolated itself from many practical solutions developed in operations research and from much of the knowledge ~~which~~ that has been gained from implementation of data base and information systems.

- A new profession of knowledge engineering is being created, without the benefit of the experience gained from complex systems analysis.

- Current expert systems lack~~s~~ flexibility and adaptability, and they are inclined to fail as they approach the periphery of these narrow domains (areas of expertise). They do not know the limits of their knowledge, and cannot adjust to the knowledge level of the human expert.

- Current experience indicates that excessive maintenance is required to ~~assure~~ ensure acceptable performance and systems quality.

D. DON'T UNDERVALUE A.I. CONCEPTS

- As information systems become more complex, it becomes essential to formalize how data and information are being used--in other words, how they are expected to contribute to human knowledge.

  - End-user access to large central data bases makes them subject to misinterpretation and to the generation of conflicting information.

  - As more information becomes available, there is a tendency for quality (or the ability to assimilate) to diminish, and the acquisition of knowledge becomes more complex.

  - Individuals using available data and information are using personal computers and human knowledge to generate more information by developing their own systems. It becomes essential to understand the "human element" applied to this new information.

- It becomes essential to apply advanced tools from operations research, statistics, and AI if decision support systems are to be effective in such a complex environment. In other words, there is tremendous uncertainty in the emerging environment.

- Intelligent systems are required--especially in the sense of being able to explain how and why they reached the conclusion (or generated the information) they did.

- As systems become more complex, it becomes increasingly difficult to make them "user friendly", intelligent systems are required to guide and educate ~~the~~ users.

E. A POSITIVE COURSE ~~TO~~ OF ACTION

- ~~Establish~~ an advanced technology research function should be established under the IS department. AI research has had a narrow focus in its application to business problems. Knowledge base development is where data base development was twenty years ago, and it is more complex. Basic knowledge concerning OR, statistics, information theory, decision theory, and library science will be required in addition to AI.

- There are few practical applications of expert systems at ~~the~~ present ~~time~~ and the narrowness of their domains dictates that great care be exercised in exploring the potential of expert systems. This selection problem is compounded by the extremely broad array of "potential applications" ~~which are~~ currently being publicized. A non-critical problem would probably be best.

- Tools such as "LISP machines" are currently dictating the problems being addressed by AI. LISP is an old language ~~which~~ that has never been formalized (there ~~are~~ is practically an infinite variety of AI languages). General claims of systems development productivity should be viewed with a high degree of skepticism.

- Management should concentrate on quality and cost. Do not apply AI in order to be on the "leading edge" especially when better solutions are available. Expert systems are going to be costly in terms of both machine and human resources. In addition, they may not work, and you may not know it.

III

III WHAT IS ARTIFICIAL INTELLIGENCE?

A. A BRIEF HISTORY

o In interpreting the work of Charles Babbage on the Analytical Engine, Ada Augusta, Countess of Lovelace, stated that:

- "It is ~~desirable~~ good to ~~guard against the possibility of~~ be wary of exaggerated ideas that ~~might arise as to~~ bound from the process of the Analytical Engine. ~~In~~ in considering any new ~~subject,~~ topic, there is ~~frequently~~ often a tendency, ~~first~~ initially to overrate ~~what we find to be already interesting or remarkable; and~~ the technology by emphasizing the interesting or amazing aspects of it. Then, when we realize it doesn't meet our expectations we tend to (Under-serve) ~~secondly, by a sort of natural reaction to~~ undervalue the true ~~state~~ condition of ~~the case, when we do discover that our notions have surpassed those that were really tenable.~~ the technology.

- "The Analytical Engine ~~has no pretensions whatsoever,~~ does not pretend to originate anything. ~~It~~ The device can ~~do whatever we know how to order it to perform.~~ perform only what it is ordered to perform. It can follow analysis; but ~~it has no power of anticipating~~ cannot foresee any analytical relations or truths. ~~Its province is to assist us in making available~~ The purpose and capacity of the device is to help make what we are already ~~acquainted with.~~ aware of."

o ~~These~~ Lady Lovelace's sage observations are as appropriate today as they were over 100 years ago. The tendency to "overrate" and then "undervalue" can clearly be seen whether we consider the "electronic brains" of the 1950s, the higher level languages in the 1960s, data base management systems in the 1970s, or

decision support systems in the 1980s. AI and expert systems will be subject to the same tendency, and if Lady Lovelace's advice can help us moderate this tendency, she deserves more recognitions than having programming language named after her.

o However, the fascination with a "thinking machine" simply will not go away despite the clearly stated lack of "pretensions" on behalf of the Analytical Engine. Alan Turing went to some length to present contrary views in his exploration of whether machines could think (computing machining and artificial intelligence). After devising the "Turing test" to resolve that question and concluding that a machine could be constructed which could pass that test, he specifically addressed Lady Lovelace's analysis of the Analytical Engine's capability--specifically the implication that a machine can never do anything really new. Turing's response to Lady Lovelace's "objection" was as follows:

- "The view that machines cannot give rise to surprises is due, I believe, to a fallacy to which philosophers and mathematicians are particularly subject. This is the assumption that as soon as a fact is presented to a mind all consequences of that fact spring into mind simultaneously with it. It is a very useful assumption under many circumstances, but one too easily forgets that it is false. A natural consequence of doing so is that one then assumes that there is no virtue in the mere working out of consequences from data and general principles."

- It should be pointed out that Turing did not see fit to even discuss the question of "can machines think?" because he considered it meaningless. he conjectured that by the end of this century one would be able to speak of machines thinking without expecting to be contradicted. He then stated that "Conjectures are of great importance since they suggest useful lines of research."

- Some of the most creative minds of this century have been engaged in such research (and speculations) since Turing established his test and the case is still being argued.

o INPUT has reviewed many of the arguments pro and con on the subject and makes ~~has~~ the following observations ~~to make:~~

- At the time that Turing made his conjecture, the art of computer programming was not very well understood. It was assumed that computers would "do what you told them to do" even if what you told them to do was "something new".

- There is no question now that machines can ~~"give rise to surprises"~~ us. Unfortunately, most of these surprises have been of the information variety; they are called bugs and they admittedly are present in programs of even moderate complexity.

- It is INPUT's ~~our~~ intuitive judgment that information surprises will be a major factor in the acceptance of all machines and systems exhibiting artificial intelligence, and especially if such machines are deemed to be thinking. It is likely that humans will be less tolerant of errors on the part of machines than they would on the part of other humans--if for no other reason than ~~the fact~~ that humans are better able to cover their mistakes.

o During the giant electronic "brain" period of the 1950s there was considerable speculation about whether computers could think, and the public and most "experts" concluded that ~~Lady Lovelace was right--~~ "Computers ~~could~~ can do only ~~do~~ what you ~~told~~ tell them to do." Thinking machine research disappeared into the universities under the name of "complex information processing" and "machine-aided cognition". Then John McCarthy coined the term "artificial intelligence" in the 1960s, and it stuck. However, even today there are those who feel the term is ~~not~~ inappropriate and might stir up the old public controversy ~~before the public~~ again. It probably will--if it already hasn't.

semantic and syntactic understanding to produce quasi-intelligent computer-to-human communication.

- Image signal interpretation in which intelligence is combined with

B. ARTIFICIAL INTELLIGENCE DEFINED

o The best way to define AI is through the research areas ~~which~~ that have received emphasis over the last twenty years. However, AI ~~it~~ has it's roots in the desire of behavioral psychologists to model human cognitive behavior, which involves hypothesizing, reasoning, perceiving, associating, speculating, learning, guessing, and drawing conclusions. These are complex processes ~~which~~ that vary significantly from individual to individual, to the degree that measurements can be made, ~~and~~ implementation in the hardware and software of the human brain remains something of a mystery ~~below~~ beyond the left brain and right brain generalization and understanding of sensory locations. (In other words, we don't have a very good storage map, much less ~~an~~ an understanding of how the program works.)

o AI has fundamentally been an effort by ~~man~~ humans to create machines in ~~his~~ their own image, at least ~~in so far as~~ to the extent of having them exhibit human cognitive behavior. Some of the major sub-fields of AI in which modeling has been done are:

- Natural language interpretation ~~where~~ in which intelligence is combined with visual imaging and pattern/object recognition to permit computer analysis of, and response to, the external world.

- Speech interpretation ~~where~~ in which intelligence is combined with aural processing and recognition to permit verbal ~~man~~ human-to-machine communications.

- Robotics ~~where~~ in which intelligence is combined with mechanical servo-mechanisms to produce "by-reason" rather than "by-rote" machine physical processing.

- Knowledge-based expert systems, where intelligence is applied to computer solution of problems whose complexity or inexactness and reasoning to solve.

o When reviewed against the interests of behavioral psychologists in human cognitive behavior (hypothesizing, reasoning, learning, etc.), it becomes apparent that the intelligence applied in most of this research areas has yet to address the primary issues except to expose the complexity of the human cognitive processes.

- Indeed, at the 1984 American Association of Artificial Intelligence convention in Austin, Texas, Marvin Minsky (one of the many "fathers of AI) specifically pointed out some of the deficiencies of AI research: "AI has been shockingly deficient in developing theories of machine learning. (For example) only a little AI research has gone into the study of machines that learn by analogy. People avoid hard and important problems just as long as they possibly can."

- Concern about AI research has also been expressed by IBM Vice President and Chief Scientist Lewis M. Branscomb, who has stated, that

  . "The extravagant statements (about AI) of the past few years have become a source of concern to many of us who have been excessive claims by researchers in other fields lead to unreasonable expectations by the public." (Shades of Lady Lovelace.)

  . Expressing his concern about limited progress of AI research in a number of areas, Branscomb pointed out that only 70 papers had been presented at the 1984 AAAI conference whereas a materials science conference of that size would have had 700 perhaps even 1400 papers.

o However, the major sub-fields of AI research do have potential for substantial benefit and should not be "undervalued" even if the tendency to "overrate" them is already apparent (at least to Mr. Branscomb). INPUT's observations concerning the selected sub-fields is as follows:

- Natural language interpretation of high quality is essential if the Turing test is to be met. (The Turing test essentially states that if a human sitting at a terminal having a dialog cannot determine whether he/she is conversing with another human or a machine then the machine can be credited with intelligence else the machine cannot be considered intelligent.) While "quasi-intelligent" conversation remains the current goal of natural language interpretation, natural language systems do not as yet have the ability to pass the test and there is considerably speculation about whether they even will, and what it would really prove any how. Tests other than Turing's "imitation game" have been suggested tests which get at the underlying understanding of the machine by having the machine explain why it is doing what it is doing or responding the way it is.

- The fact remains that conversations between humans and machines are at best strained and generally more unnatural than natural. Such a The respected computer scientist as Dr. Edgar W. Dijkstra considers anthropomorphism (as represented by the desire to converse with computers in natural language) to be unfortunate.

> "And now we have the fad of making all sorts of systems and their components "intelligent" and "smart". It often boils down to designing a woolly man-machine interface that makes the machines as unlike a computer as possible: the computer's greatest strength--the efficient embodiment of a formal system--has to be disguised at great cost. So much for antoropomorphism."

He decried the fact that humans take great pains to make computers as unlike machines as possible. Dijkstra remarked that humans disguise the greatest strength of computers by attempting to make them human like.

- It is INPUT's opinion that the quest for natural language communications between ~~man~~ humans and machine is not only costly but would prove unsatisfactory for both the ~~man~~ person and the machine.

- However, the author remembers the first time Autocoder on the IBM 705 typed out the following message on the console typewriter: "NO TAG ERRER, ASSUME TAG ERROR" in the middle of a twenty-minute asembly and continued running. The machine had "guessed" right and ~~it would~~ saved a big chunk of machine time. It makes no difference, then, whether dictionaries, punctuation programs and passing algorithms are really AI. Even limited intelligence and understanding of natural language is valuable and research should continue.

- It certainly will be possible for machines to "read", classify, cross-reference, and file documents with increased sophistication and this is going to be necessary as we proceed toward electronic offices.

- Image signal interpretation can be viewed as still being in it's infancy when it comes to playing with various shapes and colors of blocks, but substantial practical applications are already possible based on pattern recognition research. It is possible to recognize hand-written symbols sufficient to incorporate automatic reading and encoding of data from documents. Realizing that some hand-written information has traditionally be unreadible by human beings ~~(the medical profession being a glaring experience)~~, it is not beyond the realm of possibility that humans may eventually be required either to write so it is machine-readable ~~by machine~~ or to learn to use a keyboard. It is not important that machines "are" and even remember like human beings. In some cases, they probably will be better, and in others they will never be as good. Once again, continued and extended research is highly desirable, but the anthropomorphism deplored by Dijkstra remains a problem.

- Speech interpretation sufficient for many practical applications is already possible and full transcription of voice documents is certainly attractive. All of the comments concerning natural language and image signal interpretation apply here, but a few additional observations are required.

- It may be both practical and useful to use "voice prints" ~~to have individuals~~ for withdraw money from banks or unlock ~~their~~ the automobile, but ~~there~~ are a lot of speech interpretation applications ~~which would~~ seem to have limited practical value.

- For example, except in situations where hands are ~~not~~ available (during a ear or in some cases) flying an airplane, physically handicapped, etc.) there is ~~limited~~ little need to ~~talk to~~ activate a machine by voice. In most instances, ~~where~~ punching a few buttons would suffice--for example, in "dialing" a telephone or even "looking up" a number.

- ~~Then, more importantly, there is the~~ another important issue remains regarding speech interpretation: whole class of people ~~who~~ claim ~~they are~~ to be waiting to talk to computers before they will use them--it is INPUT's opinion that there are precisely the people who will never know what to say to their computers anyway.

- Robotics in its ultimate form would integrate all of the other sensory and cognitive research into a neat package ~~which~~ that would emulate a human being or perhaps improve on the model. It should be pointed out here that a great many of the concepts of robotics (of all varieties from automated machine tools to the more advanced versions) came from Norbert Wiener, who can be called the father of cybernetics, which existed prior to AI. He drew the parallel between biological and electrical systems in terms of feedback loops, which are so essential to any complex system. (Wiener remained cautious about the ~~human~~ human-machine relationship until his death, especially in supporting decision making as proposed by his friend ~~and contemporary~~ John von

Neumann~~, more on that later.)~~ Robotics have a broad range of applications ranging from practical factory automation system to the whimsical personal servant~~, and, on~~ to computers ~~which~~ that build new (and perhaps better) computers. ~~Just a few interesting examples and speculations:~~

- ~~A robot has been developed which wanders around until it's battery runs low at which point it snugges up to a power outlet and plugs itself in--it should not be any tricks to build a mobile telephone into it and be able to call it by name (Watson) when you want to use it (but do you really want to?)~~

- ~~A robot recently made the local newspapers (San Francisco) by wandering around in a restaurant, talking with customers, and passing out passes to a movie--in the process managed to disrupt the restaurant (by running into things) which probably did not permit soliciting any how (this is not considered a practical applications).~~

- ~~A "practical robot" incorporating most of the advanced AI research presented thus far would read all of the paper documents (including computer produced) coming into the office, tell you what it was doing with those it was retaining and ask if it was OK before shredding the worthless ones (including bulk mail)--but there must be a better way.~~

- ~~The point is that anthropomorphic robots have highly questionable payback (or desirability) and may be diverting AI research efforts from more practical pursuits.~~

- Knowledge based expert systems is where the real activity is these days and the next section will explore this AI sub-field in detail. However, there are two important general considerations in such systems:

- It is easier to "hide the human being in the mechanism" in such systems and pretend they are exhibiting intelligence, and you can be sure this will be done.

- It is more difficult to detect "unpleasant surprises" in an expert system than it is in say, a robot, banging around in a restaurant and it is also easier to hide mistake (either intentionally or unintentionally).

- The combination of these two factors has high potential for both hardware/software/knowledge problems with the artificial systems and for human skulduggery this is a bad combination.

o Before proceeding with a more detailed analysis of expert systems, it is necessary to review briefly some parallel developments which will have cross impact on each other.

A.I., O.R., AND C.C.

C. AI/OR/CC

o There is a common heritage of artificial intelligence (AI), operating research (OR), and command and control (CC) systems out of the United Kingdom during World War II, and many prominent persons crossed paths working on different parts of the "messy problems" associated with the more mobile warfare which was developing on land, sea, and in the air.

- Turing was working on the machines which would be used to break the German's Enigma codes.

- Wiener was working on problems associated with controlling antiaircraft fire by feedback loops for position prediction.

- Breaking the codes required new techniques for interception of both aircraft and U-boats if the real payoff was to be achieved and the term "operations research" was born.

- John von Neumann was everywhere meeting with both Turing and Wiener, and concerned with both the development of computers and the atomic bomb. Von Neumann was already the leading expert on game theory which would become increasingly important in command and control systems.

o OR became popular after World War II and its techniques were extended to a number of pragmatic problems in industrial engineering and to various operational problems in transportation communications, and other complex scheduling problems. There has been relatively little contact between OR and the more advanced aspects of AI, but now there is a call from both sides for improved communications between them. There is little question that there is significant convergence between the two general disciplines. However, many areas of AI research are proceeding without the benefit of OR experience, and most OR professionals are not acquainted with AI progress.

o John von Neumann left his mark on the nuclear strategy of the United States. War gaming and strategic planning have been based on game theory concepts and ICBM's leave no time for parlimentary debate (or even very many telephone calls). computerized decision support systems are inherent in nuclear deterrence. From there AI has filtered down to command and control systems at the battlefield level.

o It is clear that the most complex (to say nothing of crucial) expert systems have been developed by DOD. Except for a few wild rumors concerning malfunction and the impact of high stress war gaming on the human experts playing the games, the general public and business community are unaware of how well such systems work. Indeed, it is probable that a good argument can

be made that such systems will be full of "surprises" if they are even put into operation. However, one thing is certain; there comes a point when a human must trust the system enough to take action or the system becomes an impediment to decision making. It is probable that expert systems will demonstrate many of the attributes of CC systems—that is where decision theory, computer analysis and AI have been applied.

o In commenting on attendance at the 1984 AAAI conference, an observer estimated that approximately one-half of the attendees were from colleges and universities; one-third from government, industry, sellers and users of AI (with the largest group of these from DOD) and the remainder from research institutes. He then went on to say: "Still largely absent, however, are representatives of the professional, industrial and financial communities toward whom much of this activity is targeted."

IV

IV KNOWLEDGE-BASE (EXPERT SYSTEMS

A. STATUS AND TERMINOLOGY

- There are two major categories of vendors associated with expert systems: those who apply AI techniques to the development of specific application systems, and those who provide the hardware software tools to assist in the development of such systems. Since the hardware software tools support research activities and apply to all of the AI sub-fields (and have been around longer), they have substantially more revenue. The companies which claim expertise in the development of "knowledge-based" or expert systems have normally been spawned out of the universities with many of the principals continuing to hold faculty positions.

- The sudden rush of activity directed toward application of AI to business problems seems to have been spawned by a number of factors including the Japanese Fifth Generation computer effort, entrepreneurial spirit, the reported success of a few research projects, and a rush to preempt names such as Cognitive System Inc., Teknowledge, Inc., and IntelliCorp.

  - The "threat" of the Japanese Fifth Generation provided a focal point of commercial attention for AI, much of the current activity can be classified as the response of the free enterprise system.

  - However, the sudden activity in AI has also been inspired by the successful commercialization of genetic engineering and the winter

undercurrent

1985 issue of AI Magazine contains an article "Comparing Artificial Intelligence and Genetic Engineering: Commercialization Lessons." The abstract ~~states, in part:~~ remarks that "~~Artificial intelligence is rapidly leaving its academic home and moving into the marketplace. There are few precedents for an arcane academic subject becoming commercialized so rapidly. But,~~ genetic engineering, ~~which recently burst forth from academic to become the~~ the academic foundation for the ~~hot~~ new biotechnology industry, ~~provides useful insights into the sites of passage awaiting the~~ essentially sets an example for AI of the rapid commercialization of ~~artificial intelligence."~~ an academic subject. There are, otherwise, few precedents for a subject such as AI moving so quickly from academia to the marketplace.

- Then, of course, the media (and market research firms) are also involved in stirring up the pot.

  - The press was out in force at the 1984 AAAI conference and John McCarthy, during his President's Address, projected put an equation on the screen, "so that the representatives of the press won't think that they've understood everything ~~that~~ they've heard here." There (for seen)

    and remarked that it was

    This prompted some later discussions ~~to~~ on how press relations could be improved.

  - But, not to worry, the market for LISP machines alone ~~was~~ has been forecast to be $1 billion by 1990 by some market research firms — Now, with these success figures, what venture capital firm could resist the new terminology? There are some pluses in ~~publicizing~~ "spreading the word."

o The computer industry is not unacquainted with the use of new terminology, but AI and expert systems have arrived on the scene with a while new set. In order to understand expert systems, it is necessary to define some fundamental terms and concepts:

- Symbolic representatives (long) are the substituting of real-world objectives, concepts, and relationships by abstract symbols (letters, numbers, or special characters) ~~which~~ that can be manipulated logically by programmed computer processes.

- Symbolic inference is the derivation of new facts or previously undeclared knowledge from existing information through the application of formal and informal logic and search techniques.

- Heuristic search implies the use of rules of thumb for learned and/or experientially derived human expert methods used to narrow a search and guide it toward resolutions.

- Confidence factor is a numeric value associated with a fact or heuristic to indicate an individual or combination weighting of reasonableness or predictability.

- Fuzzy set is a logic formalism expressing the relative set membership of an item in which inclusion/exclusion is not absolute but is represented on a continuum of values.

wp: Continuum

- Inductive discovery is reasoning from individual observations toward a general principle.

- Deductive reasoning is reasoning from a general principle toward individual observations.

- Public knowledge is that which is generally available through text books, encyclopedia, seminars (proceedings), etc.

- Private knowledge is that which is individually developed, usually through learning, experienced, experimentation, etc. (And, through what you have probably called systems analysis ~~and is~~ now being referred to as knowledge engineering.)

o ~~Back~~ during the days when the mystery of the new "electronic brains" was being dispelled, the presient of a major transportation company concluded

NB: each instance is 'tis 'tain't

that computers were really no mystery at all as he so aptly put it, "All a computer can do is say: 'tis, tain't; 'tis, tain't..." The simplest explanation of an expert system is that it can say "maybe" and proceed from there. In other words, the human brain is not binary, it is not ~~a "'tis-tain't" machine but a~~ have a "'tis--'tain't" process, but a "maybe--I wonder" process. The expert system attempts to emulate human decision making and give expert "consultation and advice" in a well-defined and narrow domain of expertise.

o Expert systems have two major components, the knowledge base and the inference engine, but can be supported by other sub-fields of AI such as natural language interpretation, image signal interpretation, speech interpretation etc., as shown in Exhibit IV-1.

IV-1

- The knowledge base is normally described as ~~being a~~ being a loose and ~~usually~~ generally unstructured ~~collection~~ set of facts, assertions, inferences, observations, hypotheses, rules, and procedures specifically related to a particular domain of activity (such as medical diagnosis). However, it soon becomes necessary to confront the fact that some structure is required and this has resulted in the following:

  No underscore

  . There is a recognized need for expert systems to have data base access.

  . Both expert systems and data base management systems are recognized (by at least some of those involved in AI) to rely on data or knowledge representation models and they both involve some inference making.

  . However, they have evolved separately and have placed emphasis upon different aspects of the decision-making process ~~with~~ the result is that it is difficult to integrate existing expert systems (of which there are precious few) with existing data base management systems (of which there are many).

- This has and will result in new data base systems ~~being developed~~ for expert systems.

  - The Japanese Fifth Generation effort has opted for the relational model (with at least some hardware implementation) and this seems to be a wise choice (although performance of even the fifth generation may be strained).

- The inference engine manipulates and applies the knowledge bases in progressing toward a problems solution (or decision). It extracts, matches, and applies (tests) the facts, assertions, hypotheses, rules, etc. contained in the knowledge base against the statement of the problem, refining both problem statement and solution as it goes along. Fundamentally, the techniques used are those listed under the definitions of terms and concepts.

- The potential support of expert systems through other AI sub-fields is obvious when one considers the ~~required~~ human interface required for human-machine interaction, and the desirability for computer assistance in screening and structuring of public knowledge sources on specific domains. For example:

  - The human expert should be able to pose ~~his~~ problem in natural language and enter ~~his~~ own knowledge in the simplest form (pictures, charts, etc.).

  - Scanning, classifications, data information ~~reduction~~ and storage of printed material and audio-visual information to build domain knowledge bases (from both public and private knowledge) would do much to facilitate the development of expert systems.

wf: of "'tis 'n 'tain't" fame

- During the research for this study, a popular business publication listed the transportation company of "'tis 'n 'tain't" fame as being a user of expert systems. A telephone call was made to the company's IS Vice President. ~~resulted in the following:~~ The call resulted in the respondent's

  - An admission that despite a lot of talk about expert systems and AI (the respondent had recently attended a meeting at which ~~where~~ a retired Air Force General had discussed the subject), he was not really sure what made an expert system different.

  - When told that a publication had listed his company as an expert system user, he stated that they must have been referring to a dispatching system they used, but he commented: "As far as I am concerned that system is straight LP (linear ~~limited~~ programming)." (The company has an OR unit under the MIS department.) we might add

  - The respondent's ~~His~~ evaluation of the dispatching system was that it did a better job than some of the newer dispatchers but could not compete with the really experienced ones. It's main value was ~~had been~~ determined to be in training where it permitted dispatchers to be moved to new territories and become productive more rapidly.

- There are clear ~~differences~~ distinctions between the LP dispatching system and the technical definition of an expert system. However, some of the research projects on expert systems have revealed substantially the same results the expert system can produce results better than X percent of the human experts, but in many cases the Y percent of human experts who are better than the system are expected to handle the harder cases. (In other words, the system referral is the case, for example, to a specialist in the medical profession or the assignment of the experienced dispatchers to the difficult territory.) In which case, the computer-based system derives its primary benefit as a training or screening device. Using the transportation company example

EMYCIN and CADUCEUS?

- It is obvious that the selection of domains for the development of expert systems is very important, and aside from a few widely publicized systems in medical diagnosis (MYCIN and CADUCEUS) and mineral exploration (PROSPECTOR), there are currently few new applications. As an attendee at the 1984 AAAI conference commented:

> "The problem of AI seem to be within AI; very few applications of AI tools, successful or not, are reported. I hear some conjecture that the applications in place are so important that their success is being held in strictest confidence. More realistically, experience leads one to expect that when more has been accomplished, more will be reported."

B. HOW EXPERT SYSTEMS ARE DIFFERENT

- Essentially if you have a problem with a specific solution and can define an algorithm to solve it, you do not need an expert system. Expert systems are for "fzzzy" complex problems where human judgment and reasoning are normally employed (with varying degrees of success). The algorithms nature of OR problem solving would therefore rule out OR solutions as expert systems.

- In addition, there are a set of general characteristics which distinguish expert systems from conventional data processing systems which may solve rather complex problems (see Exhibit IV-2). These characteristics are highlighted in Exhibit IV-2.

  - Expert systems exhibit expert performance by solving problems which would normally required a human expert. Unlike a traditional DP application which is rigidly defined and leaves no room for probabilities or uncertainties, expert systems arrived at an acceptable solution, in a reasonable period of time, despite ill-defined data and processes.

IV-2

- Conventional data processing systems use arithmetic symbols and processes to represent data elements, files and logic paths. Whetever mathematical algebra, boolean algebra or relational algebra is used, they do not cope very well with abstract and vague concepts and relationships. Expert systems, however, employ symbolic structure to represent both concrete and vague abstract concepts, hypotheses, associations, etc.

- Data processing normally employs algorithmic processing which means that the process (including handling of errors) is predetermined; indeed, when surprises occur the program frequently aborts. Expert systems employ non-deterministic symbolic inference which means the sequencing of operations and the logic applied is dynamically determined by the system. Surprises are expected as a result of partial problem solutions and external events; the system adjusts itself and continues towards a solution.

- Conventional data processing systems normally employ single-path or an exhaustive search of all possible candidates in seeking a problem solution. Expert systems use heuristics or "rules of thumb" to narrow the search and guide it towards a workable if imperfect solution. Because of heuristic searching expert systems can address large, more complex, fuzzy and non-deterministic problems which could not be practically addressed by the exhaustive searches of traditional DP approaches.

- Expert systems have extensive domain knowledge, both public and private as their foundation. (This actually represents a shift from emphasis upon laws of reasoning and inference which resulted in the development of earlier cognitive systems, and is comparable to the early DP experience of developing general purpose management information systems without underlying data.) DP systems usually only concentrate on specific data required for actual processing even if those data are the result of an extensive data requirements study.

- When heuristics associated with a domain fail to result in a successful search, expert systems normally fall back to "weak reasoning" based on general logic to provide some answer. Conventional data processing systems do not provide answers if these exhaustive searches fail. This is possible because expert systems reformulate the human problem representation into ~~an~~ one suitable for more flexible computer resolution.

- Expert systems have ~~"self-knowledge"~~ of how they have resolved problems, and they explain what they have done and justify the recommended solutions. This self-~~answers~~ awareness is not required in algorithmic systems because ~~they~~ these systems have been "told what to do."

o Expert systems are usually applied to problems ~~which~~ that also exhibit a set of characteristics ~~which~~ that set them apart from those ~~which are~~ usually identified with algorithmic processing systems. ~~Among these characteristics are:~~

- The individual problems addressed by expert systems are usually large and complex, with multiple and possibly conflicting solutions.

- The problems are not predictable, being fuzzy and uncertain, and these solutions require symbolic inference and abstract logic.

- The problems are data intensive, requiring large amounts of factual and experiential data, and associated logic representation.

- Because of size and complexity, the problems normally can ~~only~~ be resolved only by using heuristic search; they require "expert" judgement and are not amenable to algorithmic solution.

o At this point, it seems appropriate to point out that as more data become available for both planning and control purposes, practically all problems become complex and uncertain. As personal computers are used to create

distributed and personal data bases, these may be viewed as knowledge bases in the sense that they represent a particular "way of looking at things" that an individual (or organization) uses to make decisions. In other words, as more data and information become available in the decision-making process, problems become "fuzzy" and more alternatives become available. It is obvious that the purported merits of expert systems are a logical extensions of decision support systems. The question becomes are they a practical extension?

C. HARDWARE/SOFTWARE TOOLS

1. PROBLEMS OF IMPLEMENTATION

o By definition expert systems are complex, and it should be obvious that if the problems they are attempting to solve were easy more progress would have been made. Hayes-Roth, Waterman, and Lenat in their book Building Expert Systems illustrate the components of expert systems with a general flow diagram (Exhibit IV-3). In it's ideal implementation such a system goes well beyond all of the advanced work which has been done in fourth-generation languages, data base management systems, and operating systems (the scheduling is substantially more complex).

- The language processor provides a narrow problem-oriented language specific to the expert user.

- The "blackboard" provides for recording intermediate results.

- The knowledge base contains facts, heuristic planning and problem-solving rules.

- The interpreter applies the rules.

- The scheduler controls the order of rule processing.

- The consistency enforcer adjusts previous conclusions when the knowledge (or data base) changes.

- The justifier rationalizes and explains the systems behavior.

- Considering the fact that each domain (application) is ~~specific~~ highly specialized, it is understandable ~~should be obvious (without detailed explanation)~~ why the number of expert systems developed to this point has not been large.

o The reason that expert systems are receiving so much attention is that hardware and software tools to facilitate development are beginning to appear in the marketplace. Before discussing these tools it should be pointed out that there is currently an extreme shortage of people who are knowledgeable ~~of~~ about these tools, or who have any ideas of how to develop expert systems. Indeed, the tools themselves add another level of complexity in the development of such systems. (The situation is comparable to having MVS/XA, IMS and PL/I descend upon you at one time ~~and not have~~ without any systems programmers being available.)

o Therefore, ~~while~~ even though expert systems may be the logical extension of decision support systems, available tools and personnel will determine whether they are practical.

2. ~~AI~~ LISP MACHINES

o ~~Mention AI and most people~~ When people hear AI mentioned, most of them think of LISP. As one observer at the 1984 AAAI Conference stated: "~~While~~ It may not be entirely correct that AI is what LISP makes feasible, but the coincidence of LISP with accessibility to large-scale computing power on very economical terms has provided much of the energy for the AI explosion."

- LISP is not a new language, it has been around nearly as long as FORTRAN. Unfortunately (or fortunately, depending upon how you want to look at it) LISP has not been formalized ~~but~~ It has been left to grow according to local requirements (most of which exist in the universities and research centers). ~~At the AAAI conference, David Touretsky of Carnegie-Mellon University emphasized these important characteristics of LISP:~~

There are three important characteristics of LISP:

  - The ability to evaluate words and lists, rather than numbers and tables.

  - A facility to handle truth tables and boolean algebra (the "if" then ~~else's~~ "operational concept ~~you are familiar with~~ which most people are familiar),

  - An easy going environment that ~~which~~ makes it easy to stop, take stock, and make changes ~~as you go~~ while proceeding (during development ~~during development~~).

- Of course, LISP isn't the only symbolic language (in fact, LISP itself ~~is~~ includes many languages), and the Japanese ~~has~~ have elected to support PROLOG in their fifth generation project. At the "first national meetings" of the Symbolics LISP Users Group (SLUG) a progress report stated the following:

  - "High-performance" PROLOG would not be out soon for the Symbolics 3600 (Symbolics is only one of several companies currently marketing symbolic processors).

  - PROLOG KR (which ~~had been~~ was developed in Japan), however, was immediately available.

  - In addition, LOGLISP, a combination of LISP and PROLOG developed at Syracuse University, was ~~stated to be~~ reportedly available.

  - And an ADP project was reported to be well underway ~~(wonder what the fair Lady Lovelace would say about these Symbolics 3600)~~.

For anyone having fought battles over language standardization, it is readily apparent that neither LISP nor other languages for symbolic processors have been formalized. The hope has been expressed that the commercialization will lead to formalization, but right now the users seem to be making demands for a polyglot, and the hardware vendors may be hard put to deny their wishes.

- Exhibit IV-4 lists some of the workstations hardware ~~which is~~ currently being marketed to support the expert systems development market. Personal computers microcoded with LISP and PROLOG are currently under development, and this should bring the cost down substantially. (Stet) IV-4

- However, LISP has addressed primarily the interpreter portion of the expert system development problem; and, as mentioned previously, the building of the knowledge base is now recognized as all important. Fortunately, it was recognized that the "shell" of hand-tailored expert systems could be salvaged and that has led to the emerging expert system software industry.

3. EXPERT SYSTEM SOFTWARE

- The importance of knowledge in intelligent problem solving (expert systems) is summarized in Exhibit IV-5. ~~The~~ acquiring and encoding ~~of~~ knowledge (by those scarce knowledge engineers) ~~is~~ are the most difficult and costly ~~part of~~ steps in building expert systems—just as problem analysis and data base development have been in conventional data processing systems. IV-5

- This has led to the development of a number of commercially available software shells which are beginning to appear in the marketplace. Some of these tools are listed in Exhibit IV-6. ~~While~~ Although a detailed analysis is beyond the scope of this study, some basic analysis is warranted; it is obvious from their prices that they vary substantially in capability and hardware supported. However, they can be generally categorized as follows: IV-6

- The micro-based systems (Expert Ease, M.1. Personal Consultant, and TIMM) typically process knowledge represented as IF:THEN and using backward chaining. Essentially, the system starts with a hypotheses, finds a rule whose premises (IF) supports the hypothesis, and then attempts to verify the premise by searching the knowledge base for a fact which supports the premise. If the fact isn't found, the system searches for a rule which can be used to infer the fact. The process continues until the hypothesis is either verified or disproved. They provide only limited communication with the knowledge engineer concerning flow control (the order in which the system asks questions).

- S.1, obviously a more expensive system for more powerful processors, does provide for interpretation of control terminology similar to that found in regular programming languages.

- Backward chaining works well for problems where selection from a number of possibilities is appropriate for the domain (applications). However, sometimes the application cannot be specific in that manner or the number of possible solution is too large to handle on a micro because of processing or space limitations. In this case, they are appropriate only for developing prototype systems or for training knowledge engineers.

- The more advanced commercial shells such as KEE, ART, and S.1 incorporate multiple inferencing techniques by supporting not only backward but forward reasoning. This means that you may start with a fact and find a rule whose premise is verified by the fact. Thus, the more advanced models are appropriate for more applications.

o The software development shells are attached to the knowledge base which it is being built (see Exhibit IV-3). Once the expert system (application) is completed the development engine is removed and system is turned over to

the users. The question now becomes: what types of applications are we talking about?

D. APPLICATIONS AREAS

Please answer in this order

- Ten major categories of expert systems have been identified for expert systems: These include:

(10) - Control systems which adaptively govern the overall behavior of a system. These apply to industrial production, nuclear power plant control, investment portfolio management, ~~etc.~~ and so on.

(8) - Repair systems which develop and execute plans of repair for a system. ~~For~~ example, computer hardware systems and communication networks.

(7) - Debugging systems which prescribe specific remedies for malfunctions, and are especially appropriate for computer software programs and digital circuitry. (If you have tried to determine how and why some programmers can debug so much better than others, you will understand some of the complexity of building an appropriate knowledge base for expert debugging systems.)

(9) - Instructing systems which diagnose and "debug" student behavior and enable customized computer-aided education. ~~(It is obvious that this potential market alone is so large as to defy quantification--the problem becomes the domains to be addressed and building the necessary knowledge bases.)~~

(5) - Planning systems which explore possible future actions to produce the necessary steps to accomplish a goal. Applications would include

automatic programming, project planning and remote planning. ~~(For later use of the debugging systems?)~~

6 - Monitoring systems which compare observations against normalized criteria to determine deviations, ~~for~~ examples are financial management monitoring, air traffic control, and nuclear power plant monitoring.

4 - Design systems which configure objects or processes under goal constraints and can be used for computer configuring, financial budgeting, and circuit layout.

3 - Diagnosis systems which infer malfunctions from observables and include not only the classic medical diagnosis but also electromechanical and software failure analysis.

2 - Predicting systems which infer likely consequences from current observations and historical trends. The classic weather forecasting problem falls in this area ~~plus~~ along with crop estimation and battlefield movement prediction~~, and, theoretically, price-performance and timing of future IBM product announcements, with impact on residual values.~~

1 - Interpretive systems which infer situations from observables and can be used for many of the expert systems support functions (speech understanding, image analysis, signal interpretation, etc.) as well as a variety of intelligence operations.

o ~~The categories of expert systems categories are normally arranged in a more logical manner than INPUT adopted--this was intentional.~~ The potential applications areas ~~when arranged in a logical manner~~ lead to the conclusion that practically all information systems components (hardware, communications, data software) and paper-based information flow in support of the decision-making process) can be subject to replacement with expert systems, and with impact on the human experts (once they have contributed their knowledge), as shown in Exhibit IV-7.

IV-7

- The impact of IS personnel ~~impacted~~ is understandable--after all, knowledge-based systems are the logical extension of data base and information systems. However, the heavy concentration of potential impact among financial personnel might pose some funding problems. Of course, our example is exaggerated and arguments will be made ~~heard~~ that expert systems will merely permit highly skilled personnel to do a better job. However, the threat is there, especially for information systems personnel who choose to ignore the technological advances that ~~which~~ are possible in the systems development process. ~~The corporate controller may be reluctant to invest in an expert system aimed directly at the financial area, but the IS area is certainly fair game.~~

- INPUT doesn't think Lady Lovelace would have viewed expert systems much differently than she did the Analytical Engine. ~~Just as~~ She intuitively knew that programmers would be required, and, likewise, one can almost hear her say: "My, expert systems ~~shall~~ will certainly require a lot of people to advise them on what they ~~might~~ may do. I think I will call them knowledge engineers."

WP:

, "

Engineers.

V INPUT'S ANALYSIS

~~A. WHERE THINGS REALLY STAND~~

A SIMPLE CASE STUDY

- Perhaps the best example of a practical AI application ~~which~~ that has been developed in a university environment and actually transferred to an industrial environment is XCON. XCON is an expert system ~~which~~ that expands and validates orders for DEC computer systems and then configures the placement of components within equipment cabinets. The original version of XCON, called R1, was developed at Carnegie-Mellon University in 1979 and delivered to DEC in 1980. A brief description of this project is contained in "R1 And Beyond: AI Technology Transfer At DEC," AI Magazine, winter 1985.

- The article makes several important points ~~which are important as~~ that relate to expert systems ~~begin to emerge~~ from the research laboratories:

  - The systems as delivered to DEC in 1980 contained approximately 750 rules ~~and while~~ Although it could configure many of the orders, it required "extension and refinement" before it became dependable enough for use. ~~The article states: "DEC learned that it could not simply use the university designed system as designed in the university and keep it running" without making some changes.~~ ~~DEC discovered that~~ keep

These engineers, trained in expert systems programming and knowledge acquisition, corrected and extended the system's knowledge.

- The ~~"system's knowledge had to be converted and extended by people trained in knowledge acquisition and expert systems programming", and~~ a group was formed at DEC to assume responsibility for maintaining XCON. This group was composed of DEC engineers who were not selected for their AI expertise (although some at least ~~had~~ were familiar~~ity~~ with AI). ~~with some help~~ After ~~"~~about a year, and ~~much~~ with help from Carnegie-Mellon, the group was able to take over maintenance of XCON (which had grown to contain approximately 1,000 rules). This transfer was described as ~~"~~successful but difficult~~experience"~~.

- ~~It was stated that "DEC learned the hard way that maintaining an expert system requires specially trained personnel."~~ The system has now grown to 2,500 rules, and even though it was a difficult experience, an industrial user has been able to ~~keep~~ help complement and maintain an AI system.

- Since 1982, DEC has expanded its AI activities and has developed a number of expert systems including XSEL, which is designed to aid sales~~men~~people in tailoring systems to meet customer's particular needs. The experience with XSEL paralleled that of XCON: development at Carnegie-Mellon, delivery to DEC, and establish~~ment of~~ing a highly skilled maintenance group ~~(the XCON group "did not have enough time to maintain a separate system".)~~

- Based on the XCON-XSEL experience, DEC has established a formal training program for knowledge engineers, and a research group (Knowledge Engineering Advanced Development Group) to "~~decrease~~ease DEC's ~~dependence~~reliance on ~~university~~academic researchers".

- ~~The concluding statement of Stephen Polit, the author of the AI Magazine article seems sound: "As AI's potential has been understood and the technology practiced,~~ expectations about ~~its~~ AI's capabilities have ~~grown. People's feelings about AI's influence on the world and their~~ increased as people have begun to understand the potential of AI, and as the technology has been employed. In addition, the way people feel about the influence of AI has changed.

- The author of the article, Stephen Polit, warns about exaggerating the capabilities of AI -- in order to avoid unrealistic expectations -- and urges IS managers to reevaluate goals. He reminds readers that even new ideas and technological developments cannot solve every problem.

~~own particular organizations have grown as well.~~ AI is ~~actually~~ ~~changing the way~~ DEC and ~~many~~ other organizations ~~do business~~. ~~My only caution is to avoid inflated expectations about the capabilities of artificial intelligence and be willing to reevaluate your goals. Even revolutionary ideas and technologies do not solve every problem.~~

- really affecting the business practices of

o INPUT has several observations concerning the XCON case study, none of which ~~is~~ are intended to diminish the significance of either the achievement or the lessons learned.

- made
- notable

- XCON is essentially the replacement of a group of "technical editors" who represented a bottleneck in the order entry process. Evidently DEC had unsuccessfully attempted to "automate" this function ~~without success~~ before turning it over to Carnegie-Mellon as a research project. (DEC has had a "long ~~and~~ historical relationship with AI.")

- The fact that Carnegie-Mellon was able to develop an expert system to solve the problem does not mean that it is either necessary or even desirable to have an expert system perform configuration control. ~~For example,~~ IBM (for example) exercises configuration control in the order entry process on a more complex product line, with higher volumes, without resorting to ~~the~~ expert systems.

- ~~In other words, DEC's attempts to automate the configuration control and/or order entry process may have reflected failure on the part of DEC's internal system development function (for whatever reason). In fact, there were external indications as late as 1983 that DEC was experiencing difficulty with its internal information systems. A company executive attributed poor company performance to the fact that they "nearly lost administrative control of the company".~~

- It is not at all certain that the resulting expert system is a cost-effective solution to what appears to be essentially a data processing

problem. And, it is probable that no cost analysis has ever been made, ~~or at this point, is even possible~~ of the development costs of the system. At this point, such a cost analysis may not even be possible.

- However, for anyone acquainted with the problems of life-cycle costs, the heavy emphasis upon maintenance associated with both the XCON and XSEL projects is indeed ominous. To maintain the knowledge base, the organization needs ~~The requirement for~~ highly skilled technical personnel, having ~~with~~ a rare ~~scarce~~ or nonexistant specialty (AI). This ~~to maintain the knowledge base~~ is an extremely important lesson for anyone who would benefit from the DEC experience.

- It appears that knowledge engineers are destined to learn what IS management already knows--maintenance represents over 60% of total systems cost over the systems life cycle. The question becomes one of whether the cost of knowledge base maintenance will compound an already serious problem; and, indeed, whether these complex expert systems will even be maintainable as they grow. The XCON case study points out a specific weakness of expert systems --a weakness that ~~which~~ is already recognized in the industry--It is what INPUT prefers to describe as "system fragility".

2. EXPERT SYSTEMS FRAGILITY

- Current expert systems are frequently referred to as being "brittle" because they tend to fail on problems ~~which~~ that are not well centered within their specific domains. This is in contrast to robust solutions ~~which~~ that seem to function well at, or even beyond, the periphery of their intended use. For example, queueing networks, which were developed as a tool of operations research and industrial engineering, have been found to yield good results in complex scheduling problems found in computer operating systems--much to the surprise of computer scientists.

- INPUT prefers the term fragility ~~for describing expert systems because it is~~ as a reminder that expert systems are not robust. unfortunately they are not obviously fragile ~~unfortunately~~. There is a ~~our concern that, unlike an installed pane of glass, which is obviously brittle and even more obvious when it is broken, there with be~~ tendency to package expert systems. ~~This~~ which will result in a certain amount of protection for the contents, but will also obscure the fact that they may break without warning and be worthless. ~~Attempts to make it more robust by~~ systems packaging may ~~only~~ also result in hiding the fact that it has become worthless. The least one can do is stamp the contents: "Fragile--Handle With Care".

- This general labeling of expert systems as fragile is not capricious. Responsible researchers in the AI field are candid about the current state of the art and, at this point, the development of expert systems remains art despite the knowledge engineering label. Let us consider some of the acknowledged problems associated with current expert systems.

  - The current problems associated with brittleness center around the inability of expert systems to reason about either their own knowledge or that of others. They ~~do not know~~ are unaware of what they know or, more importantly, what they do not know. They cannot modify their behavior based upon what the user already knows about the subject. ~~(Perhaps it is an oversimplification,~~ but imagine ~~yourself~~ dealing with a human expert ~~who had~~ with such characteristics.) This may, of course, be oversimplifying the situation.

  - These defects lead to narrow and restrictive domains for expert systems because they lack flexibility and adaptability. When confronted with problems on the periphery of its knowledge base, an expert system is likely to wander off aimlessly searching for a solution because it "doesn't know what it doesn't know"; or, as an alternative, may ~~"non-gracefully"~~ ungracefully complete it's processing when it cannot ~~fit its~~ apply any of its solutions to the specific problem.

  - As a result of narrow and restrictive domains, there is a tendency to develop individual expert systems ~~where~~ in which there are substantial overlaps

in the knowledge base. For example, there is certainly substantial knowledge base overlap in the XCON-XSEL case study, and yet it was determined that two separate groups of knowledge engineers were required--one for the configuration control system and one for the sales-support configuration system. ~~(If you feel there is some similarity between "files" being associated with programs prior to the development of data base systems, you are correct.)~~

- Even with expert systems being limited to narrow domains, the human-machine dialogues are so formalized and restricted that a knowledge engineer is frequently required to interface between the human expert and the expert system in order to capture and represent knowledge. Since a knowledge-based system is by definition a feedback-loop-oriented system, it is probable that the development engine ~~(Exhibit IV-3)~~ will remain a part of the expert system long after it is in productive use and perhaps over the systems life cycle. In the simplest possible terms, the current defects of expert systems indicate a high probability of heavy maintenance by highly skilled personnel.

- The basic problem is that expert systems don't learn very well, and learning is one of the "hard problems" Marvin Minsky chastised the attendees at the 1984 AAAI conference for not addressing. ~~The~~ research in human learning has failed to reveal how it takes place in the brain ~~much less~~ the mysteries of the human mind. It is INPUT's opinion that current expert systems ~~model~~ libraries (complete with human libraries in the form of knowledge engineers) much more closely than they do the human learning process.

- A step beyond learning is decision making ~~where despite the artificial~~ intelligence label, ~~expert systems~~ adopt what is essentially an algorithmic approach. It is only necessary to examine knowledge representation in expert systems to determine that this is so; and, to the degree that expert systems are to be used for decision making, it is

important to understand the limitations of algorithms approaches to decision making in the general business environment. Knowledge representation in expert systems is currently one of the primary research areas; a brief summary of the various representation schemes ~~is as~~ follows:

- Predicate calculus evaluates arguments and their relationships to a truth table (either TRUE or FALSE).

- Rules ~~are~~ have the classic "IF:THEN" scheme of representing knowledge for purposes of inference.

- A Semantic Network is a knowledge representation for describing properties and relationships through the modeling of objects as graph nodes and of their relationships as labeled arcs connecting the ~~related~~ nodes. (It is not clear that the human brain works in this fashion, although there seems little question that an analog representation is more likely than the TRUE-FALSE, IF:THEN schemes.)

- A Frame is a knowledge representation in which descriptive, relational and procedural information about an object (or a concept) is grouped into a data structure consisting of various "slots" in which a pyramidal inheritance is inferred. (~~It~~ Although this may seem unclear, it has improved since the term was introduced by Minsky in 1975 ~~and~~ later research concluded "it is not at all clear now what frames are, or were even intended to be.")

- A Procedural is a knowledge representation in which the execution of a procedure (usually algorithmic) permits the inference of some fact, concept, hypothesis, etc. (In other words, if a procedure sorts on a certain key it implies knowledge concerning the relationship of the key elements.)

- In order to understand the fragility of expert systems, it is important to at least identify major areas of ignorance ~~which~~ that contribute to this fragility. In other words, even if expert systems ~~do not know~~ are unaware of what they don't know, it is desirable to understand what the developers (or the proponents) of such systems don't know.

A. IGNORANCE AND EXPERT SYSTEMS

- The areas of ignorance associated with expert systems are so vast and complex that they can only be touched upon in this report. We have isolated three general areas ~~which~~ that should not be ignored in considering the practical applications of expert systems: These are learning, decision making, and computer hardware performance. An appropriate, or single source for identifying these areas ~~was selected~~ is The Encyclopedia of Ignorance, Pergamon Press Ltd., 1978. The importance of understanding ignorance is succinctly presented in an essay on evolution which quotes from The Unwritten Comedy, by Phenella.

  - "To be ignorant of many things is expected."

  - "To know you are ignorant of many things is the beginning of wisdom."

  - "To know a category of things of which you are ignorant is the beginning of learning."

  - "To know the details of that category of things of which you were ignorant is to no longer be ignorant."

- ~~While~~ Although the above was written to describe a human being, it could as well apply to an expert system, and perhaps provide a rough measure of progress toward knowledgeable systems. The Encyclopedia of Ignorance represents the efforts of some of the world's leading scientists, including Nobel laureates, to categorize areas of ignorance in such diverse fields ~~such~~ as space, mathematics,

The ~~main~~ authors identify specific categories of ~~ignorance associated~~ ignorance that are important ~~of importance~~ in physics and biology. ~~Specific categories of ignorance of importance in the areas isolated above are as follows:~~ ~~with ignorance of importance to~~ There are 2 follows: ~~with the areas of expert systems~~ when considering the development of expert systems.

- In the general area of learning, ~~three separate articles pointed out the following:~~ it can be stated that:

  - Learning is dependent upon memory, and we do not understand how memories are stored in the brain. In fact, the more research that ~~which~~ has been done, the more complex the problem seems to become. ~~(Learning and Memory and the Nervous System, H.A. Buchtel and G. Berlucchi.)~~

  - We do not currently understand consciousness (or how the brain and mind relate) or whether it is of any use to us; specifically, we are ignorant of it's potential role in intelligent computer systems; 1) ~~would~~ such systems would benefit from consciousness, 2) if we could ~~they~~ be provided with it, and 3) ~~would we~~ (humans) would recognize it if they had it. ~~("Consciousness", Richard L. Gregory, formerly Chairman of the Department of Machine Intelligence and Perception at the University of Edinburgh)~~

  - Despite progress, including that associated with artificial intelligence, we ~~still~~ are ignorant of the languages of the brain. We are aware, however, ~~except~~ that it is improbable that ~~they are~~ brain languages consist of symbols occurring in a single serial string, and that is what symbolic processing is all about. ~~(The Language of the Brain, H.B. Barlow)~~

- John von Neumann was especially interested in decision making in economics and, along with Morgenstern, published Theory of Games and Economic Behavior in 1944. Some of the fundamental problems identified with the application of mathematics to the social sciences were identified at that time and, despite progress (the last Nobel Prize in Economics was awarded to a mathematician), modeling in the social

sciences remains an area of ignorance. ~~(Mathematics in the Social Sciences, C.W. Kilmister)~~

- In 1943, von Neumann stated: "Mathematics ~~has actually been used in economic theory,~~ perhaps even in an exaggerated ~~manner~~ way, perhaps, and also ~~. In any case, its use has~~ not ~~been highly~~ very successful." ly. the only change today has been the provision of quantitative data, and von Neumann's remark has had little effect.

- In observing that social phenomena are at least as complex as those of physics, von Neumann concluded that ~~it is "to be expected or feared, that~~ mathematical discoveries ~~of a stature~~ comparable ~~[illegible]~~ in ~~[illegible]~~ magnitude to ~~that of~~ calculus ~~will~~ may be ~~needed in order~~ necessary to progress in the field".

- ~~The author concludes that~~ There remains a great area of confusion (ignorance) in distinguishing ~~between~~ mathematical methods appropriate for modeling physical (mechanical) phenomena and their incorporation in ~~ontological systems (such as economics).~~ systems using game theory in the social sciences.

- Essentially, the social sciences deal with artificial systems--those involving interaction with human beings--and these are precisely what expert systems are all about. ~~In an article on Human Thought and Action as an Ingredient of System Behavior, M.M. Lehman reaches some conclusions which are especially pertinent.~~

Insert attached. next p.

- By applying Godel's theorem, he concludes: "Ignorance about the total behavior of artificial systems is intrinsic to their very existence."

- In establishing an "uncertainty principle" for artificial systems (comparable to that of quantum physics), he states: "in the case

→ • It can be proven mathematically, that artificial systems are unpredictable in their total behavior, ~~and~~ especially in regard to the actions that human beings will take.

→ • In addition, ~~it is concluded~~ it is impossible to determine the state of the total system, even by applying ~~probability~~ probabilistic judgments as in Game Theory.

of artificial system on the other hand one cannot (at the present time) predict with precision, what will be observed, how it will be interpreted, what action will be taken as a result.

- In addition, "one cannot even determine the state of the total system with any certainty by applying probabilistic judgements as in Game Theory."

- And, "one cannot even identify with absolute certainty, the totality of system measures that are involved in any given measuremnt activity. Far less can one determine an absolute and general limit to the accuracy with which even the known measures can be observed".

- Hardware performance of even conventional data processing systems has been of concern to INPUT, and we have frequently quoted ~~from~~ Hans Bremermann's article ~~on~~ <u>Complexity and Transcomputability</u>. However, the points made ~~can~~ bear repe~~tition as~~ating since they apply to expert systems.

NO UNDERLINE

- Transcomputability is described as ~~being~~ the point beyond ~~one a~~ which the computational cost exceeds all bounds that govern the physical implementation of the algorithm.

- It can be shown that many of the algorithms of artificial intelligence and operations research are transcomputable--in fact, any algorithm whose computational cost grows exponentially with a size parameter N is transcomputational for all but the first few integers of N.

WP: n

- ~~It is pointed out that~~ transcomputability applies not only to existing computers but to any ~~which~~ that can even be built (including the Fifth Generation system).

A 7 Controversy ~~concerning~~ Decision Theory and Game Theory

- Bremerman concludes that it ~~"This~~ is a disturbing ~~thought~~ to imagine that transcomputability applies to all computers, and many people ~~have~~ chosen to ignore ~~it.~~ this fact.

o In summary, knowledge-based systems are being developed while we are still ignorant in many important areas ~~which would~~ that seem ~~to be~~ necessary for ~~the improvement of~~ correcting their systemic current weaknesses (flexibility, adaptability, cost and even practicality). ~~It is hoped that by categorizing these areas of ignorance we have all learned something about expert systems. Consider, the weaknesses.~~ It is important for approved expert systems with understanding of these weaknesses.

o It is well beyond the scope of this study to analyze decision and game theory in any detail, ~~but~~ expert systems, as ~~an~~ extension of decision support systems, will have a propensity to head in that direction. A few paraphrased quotes will ~~be~~ help~~ful~~ ~~in~~ to identifying certain concerns, ~~that INPUT shares with many of those most closely associated~~ ~~juri~~ ble continuing controversy on the subject.

- "The indications are that...the best that mechanization will do for a long time is to supply mechanical aids for decision-making while the process itself must remain human. The human intellect has many qualities for which no automatic approximation exits. The kind of logic involved, usually described by the word "intuitive", is such that we do not even have a decent description of it." John von Neumann in speech in 1955.

Insert Attached

- "~~He~~ (von Neumann) thus tied the limitations of decision theory to those of automata theory, and he described the limits of devising models of the human brain in precisely similar terms. Other practioners of applied game theory, however, do not share von Neumann's caution in this regard." Steve J. Heims in John von Neumann and Norbert Wiener, 1980.

- "What applications of the theory of games do, is to reinforce the players' acceptance of the rules and competitive premises, and therefore make it more and more difficult for the players to conceive that there might be other ways of meeting and dealing with each

In 1955 John von Neumann warned that the best computers could expect to do in the foreseeable future was to provide mechanical aids for decision-making while the process would of necessity remain human. He pointed out that no automatic approximation existed for the human intellect, and that it was impossible to describe the logic involved in "intuitive" decision-making.

In a 1980 biography of von Neumann (John von Neumann and Norbert Wiener), Steve J. Heims pointed out that von Neumann had thus limited decision theory by the inability to develop computer models of the human brain. Heims then noted that others had not been as cautious in applying game theory to human decision making.

Gregory Bateson, who shared Norbert Wiener's concern about the application of game theory (especially to military strategy), felt that the mere use of game theory rules and premises

would change the "players'" perceptions of how they dealt with each other, and that they would become unable to relate on a more human basis. He anticipated that in the long term such changes would be in a "paranoidal direction and odious."

— Heims then reaches the conclusion that most large organizations (corporations and governments) are vulnerable to seeking simple, automatic solutions to problems ~~which~~ that require decisions — that there is a natural tendency to diffuse responsibility. He then warns that miscalculation ~~on~~ of the effectiveness of ~~pseudo~~ pseudoscientific solutions to complex problems is extremely dangerous.

— As INPUT concluded in Impact of Office Systems on Productivity ~~November~~ (1983): "In summary, we conclude that significant improvement in white-collar productivity depends upon a complex hierarchy of interrelated systems, all of which are approaching current limitations of mathematical science and knowledge in a number of technical fields. Information systems organizations are going to have to deal with questions of artificial intelligence whether or not the questions can be answered or the

problems solved." Good luck!

other...The theory may be "static" within itself, but its use propagates changes, and I suspect that the long term changes so p[illegible] a paranoidal direction and odious." Gregory Bateson i[illegible] Norbert Wiener in 1952.

paraphrase ok

- "But suppose for a moment that the technical limitatio[illegible] can be overcome and the summa bonum of a perfect "ga[illegible] machine" could be constructed. One would only need to [illegible] representative individual who could spell out the utility f[illegible] expectations of the particular "player-entity" (be it a government, a corporation, or a particular department). Technicians could elaborate all the strategic options, translate them into suitable quantitative terms, and feed all the information into the machine, which would make the decision according to instructions. What pattern of decision making could be more congenial to a bureaucratic organization? The decision-making process would become completely depersonalized. Moreover, if one takes the language of the theorists literally, the machine will arrive at the "optimum" decision. No organization could ask for more. As Max Weber has pointed out, institutional structures enhance their stability by becoming independent of particular individuals, and consequently formal, mechanical processes for decision making are an asset for an institution that seeks long life. In this way personal responsibility for consequential decisions is also neatly diffused, if not avoided altogether. As in the case of those elixirs of immortability that killed many a Chinese emperors up to the eleventh century, the efficacious instrument may be highly desirable, but to mistake pseudoscience for science is highly dangerous." Steve J. Heims in John von Neumann and Norbert Wiener.

B. RECOMMENDATIONS

1. ~~A.~~ WHAT IS GOING TO HAPPEN?

o ~~As this report was being completed,~~ Recently, a local Silicon Valley newspaper--the Times Tribune, Palo Alto, ~~California 2/25/85 had~~ featured an article ~~headlined~~ entitled

Market W ~~Feb.~~ Feb. 25, 1985

"Investors Losing Patience in the AI Software ~~Market~~." So perhaps this section should ~~be headed~~ carry ~~a~~ subhead "what is happening." The article pointed out the following:

- Few AI-oriented software companies ~~which were~~ started in the last few years have made many, and most have not produced a marketable product.

- Continuing needs for cash are meeting ~~resistance~~ resistance from the venture capital community.

- It is shakeout time already, and AI itself seems to be getting a bad name (before anybody really understands the products very well).

- There is a trend away from the use of the term "AI" in favor of such vague terms as "smart" and "enhanced".

- The reasons given for the problems of AI software companies ~~were~~ are not unusual—technical problems in development and conflicting ideas about the market.

o INPUT ~~does not feel~~ (sic) that AI—or, specifically, expert systems—~~and specifically, expert systems~~ will be forced back into the universities and research centers for another 20 or 30 years because of lack of funding. However, Lady Lovelace was right—we do tend ~~there is a tendency~~ to "overrate what we find ~~already~~ interesting or remarkable" and then to "undervalue the device's condition. ~~true state of the case~~". This report notwithstanding, INPUT believes we are still in the overrating stage of the process as far as expert systems are concerned, and that this will continue until a number of significant systems failures occur (as opposed to vendor failures). During the overrating period:

- A lot of tools for the development of expert systems will be sold (LISP machines ~~Machines~~ and development shells). Gradually the sales of such tools

will go beyond the universities and research centers and find their way into major software development companies and even enterprising end users.

- The software companies ~~which~~ that attempt to develop general purpose expert systems will continue to experience both technical and marketing problems.

- ~~There will be~~ a number of major expert systems efforts ~~which~~ will fail--some with a bang and some with a whimper.

- Expert systems will be developed ~~where 95% of the~~ for problems that could be solved with other approaches (operations research, data base systems, ~~through~~ systems analysis, etc.) more effectively and economically. more than 95% of the time

- Expert systems ~~which~~ that do become operational will require excessive resource use in terms of machine time and personnel costs for maintenance.

- Expert systems (or reasonable facsimiles thereof) will be blamed for failures in the decision-making process (regardless of whose fault it was).

- The IS department (or knowledge-engineering department) ~~which~~ that developed (or installed) the expert systems will be blamed for the failure.

o Expert systems will then become undervalued by everyone involved, and another ~~Lady~~ Lovelace cycle will have been completed.

o If all of this seems familiar--it should. ~~You have been there before with~~ The same cycle exists for other magic solutions to complex problems.

WHAT YOU SHOULD NOT DO

- Do not ride the overrate-undervalue rollercoaster this time~~-~~ it is becoming more dangerous as the "solutions" ~~get~~ involve higher levels of management ~~involved~~ in artificial systems (~~with~~ having the potential weaknesses INPUT has described).

- Do not remain ignorant of what expert systems can do--either for you or to you.

WHAT YOU SHOULD DO

- This report has identified several categories of ignorance ~~which~~ that underlie systems~~--~~. This is the ~~beginning of~~ first step in learning. Managers are urged to become familiar with the details of the categories of things ~~of which you~~ they are ignorant--that is, to no longer be ignorant. Essentially, this is to understand the fragility of expert systems so ~~you~~ they can be handle~~d~~ ~~them~~ with care. However, ~~you~~ this will also ~~have becoming~~ make the manager an expert on expert systems~~,~~ that is the paradox of understanding ignorance.

- ~~While~~ it is important to understand the fragility of expert systems, and this report has emphasized weaknesses to take the peak off the overrate phase of the ~~Lad~~ Lovelace's cycle. But it is also important to obtain knowledge about expert systems and their use. it is recommended that responsibility for acquiring this knowledge be centralized in a separate ~~operating~~ research function (either a single person or an entire organization) with responsibility for the following areas:

  - Operations research and statistics.

  - Information theory.

  - Decision theory.

  - Library science.

- Artificial intelligence (including expert systems).

o If there is already an operations research group in your organization, it should be organizationally located (or relocated) within the IS department and assigned responsibility for the other areas. Staffing of the other areas will depend upon the particular requirements of ~~your~~ the company. However, it is not deemed essential to rush out and recruit knowledge engineers (especially since they aren't available). The purpose of this group is to become knowledgeable in order to enhance and extend information systems to more complex problems--not to sell expert systems.

o The organization~~s~~ should be viewed as an applied research group to work with data base administration, information centers, development centers, and libraries, as well as with IS and user development efforts. Ideally, the funding for research in advanced techniques would come from consulting fees (or charge backs) from other areas. (DEC had problems with this because of high demand for knowledge engineers and had to isolate a separate AI research function, but the broader scope of this recommendation should alleviate some of the problems.)

o It would seem that information and development centers are the ideal vehicles for promoting the use of advanced techniques, for identifying promising applications areas, and for exercizing quality control of systems being developed in a distributed systems development environment (DSD) ~~(cite U-SSP).~~

- Initial expert systems ~~which will be~~ available on a commercial basis will be addressing the financial and planning functions. These are precisely the types of applications ~~which~~ that are addressed by information centers (and the DSD environment in general). Knowledge engineers (systems analysts) from the research organization should begin to explore how the information (reports) from the DSD environment is being used to support the decision-making process, and whether expert

systems (or other advanced analysis tools) can be applied. Then it can be determined whether commercially available AI products are of value.

- The tools of AI (LISP machines and software shells) are being promoted for productivity improvement. The research group should evaluate whether such tools are appropriate for either information or development center use.

- New Opportunities for Software Productivity Improvement detailed the quality control problems associated with the DSD environment. All of these problems will be accentuated with the use of knowledge-based systems. For example, one of the problems isolated in the DSD environment was "conflicting reports to management"; imagine the problems if "conflicting decisions (solutions)" become available from overlapping expert systems.

o In fact, the strongest recommendation of all involves quality control of vital data/information/knowledge. Read New Opportunities for Software Productivity Improvement and make your applied research function an extension of systems quality control; many of the advanced tools and techniques are required for quality assurance.

o A good argument can be made that as more data/information/knowledge becomes immediately available at the desk top, the choices (decisions) necessary will overwhelm the human component of artificial systems. Expert systems may be the only answer at the point of man-machine interface in order to narrow the choices; and the development, installation, and management of such systems is probably going to be your responsibility. The last advice is quite simple: "Handle With Care"--not only because the contents may be fragile, but because the package may contain a bomb!

A-SAI

ATTITUDE TOWARD NEW TECHNOLOGIES ~~DEVELOPMENTS~~

~~CERTAINTY AND UNCERTAINTY~~

The ~~Lady~~ Lovelace's Cycle--New Technology Is:

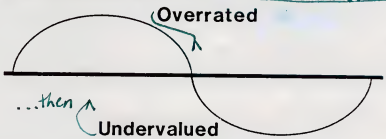

The Solution, ~~Solution, Solution.~~ to Uncertainty Is:

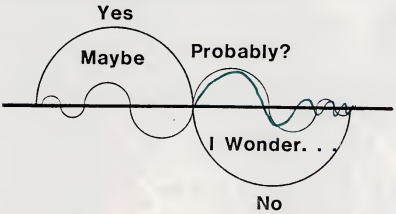

fix if possible

©1985 by INPUT. Reproduction Prohibited.

EXHIBIT II-2

INPUT®

DATA, INFORMATION, AND KNOWLEDGE

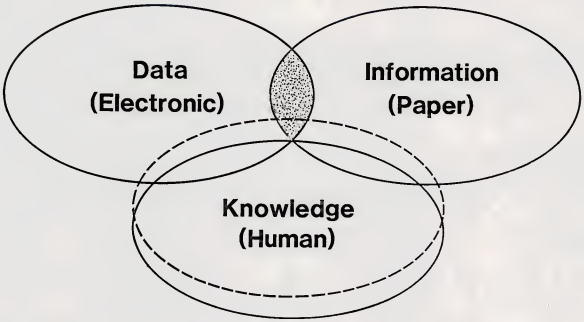

©1985 by INPUT. Reproduction Prohibited.

INPUT®

DON'T UNDERRATE PROBLEMS

- Fundamental Ignorance About Human Knowledge Acquisition and ~~[illegible]~~ Memory
- Poor Interdisciplanary Communication's
- Inexperienced Personnel
- Narrow Domains (Applications) of Expert Systems
- Poor Quality of Systems

©1985 by INPUT. Reproduction Prohibited.

EXHIBIT II-4

INPUT®

DON'T UNDERVALUE A.I. CONCEPTS

- I.S. Complexity Requires Formalization of Knowledge
  - Large Data Bases
  - Information Explosion ~~Results~~
  - Distributed Systems Development
- New Tools are Needed
- Intelligent Systems Are Required:
  - To Justify Conclusions
- ~~Education~~ Intelligent Systems Are Needed To Educate Users

©1985 by INPUT. Reproduction Prohibited.

EXHIBIT II-5

INPUT®

A POSITIVE COURSE OF ACTION

- Establish a Research Function
  - O.R.
  - Statistics
  - A.I.
  - Information Theory
  - Decision Theory
  - Library Science
- Select Expert Systems Applications Carefully
- Evaluate A.I. Tools
- Concentrate on Cost and Quality

©1985 by INPUT. Reproduction Prohibited.

EXHIBIT IV-1

AN EXPERT SYSTEM SUPPORTED BY AI

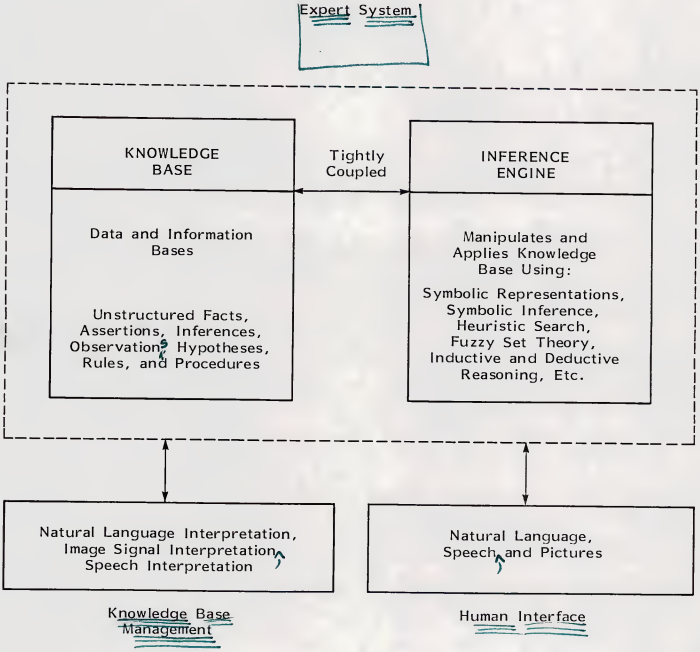

©1985 by INPUT. Reproduction Prohibited.

EXHIBIT IV-2

DISTINGUISHING CHARACTERISTICS OF EXPERT SYSTEMS

- They Exhibit Display Expert Performance
- Symbolic Representation Employ
- Symbolic Inference Employ
- Heuristic Search Use
- Extensive Domain Knowledge Have Foundation
- Fall Back to "Weak Reasoning" for Answers
- Problem Reformulation
- Self-Knowledge Have

©1985 by INPUT. Reproduction Prohibited.

EXHIBIT IV-3

THE COMPONENTS OF AN EXPERT SYSTEM

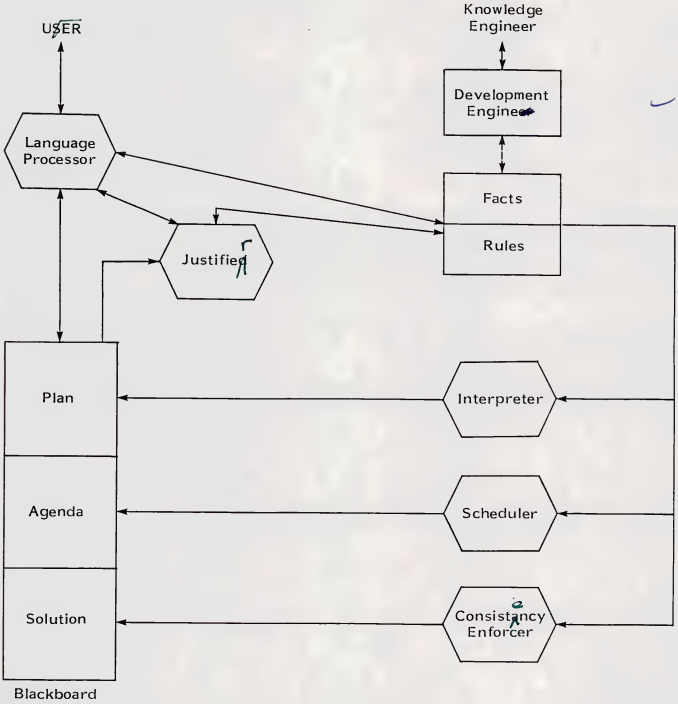

©1985 by INPUT. Reproduction Prohibited.

EXHIBIT IV-4

EXPERT SYSTEM HARDWARE

VENDOR	EQUIPMENT	PRICE
Lisp Machine, Inc.	Lambda Lambda 4 X 4/+	$66,500 $45,000
Symbolics, Inc.	3600	$59,000 - $84,000
Xerox, Inc.	1108 (Dandelion) 1132 (Dorado) 1100 (Dolphin)	$25,000 - $50,000 $150,000
DEC	VAX	
Tektronix, Inc.	4404 Artificial Intelligence System	$14,950
Texas Instruments	Explorer	$52,000 and up
Perq Systems Corp.	Perq AI Workstation	$40,000

©1985 by INPUT. Reproduction Prohibited.

EXHIBIT IV-5

IMPORTANCE OF INTELLIGENT PROBLEM SOLVING

- Knowledge = Facts plus Beliefs plus Heuristics
- Success = Finding a practical solution - one ~~whig~~ that is good enough and is not too expensive
- Search efficiency directly affects success
- Aids to Efficiency:
  - High Quality Knowledge Base
  - Elimination of Blind Alleys and redundant ~~competition~~ utation
  - Improved Speed of Operations
  - Coordination of Knowledge ~~Services~~ success
  - Various Levels of Reasoning
- Contribution~~s~~ or to Problem Difficulties
  - Bad Data or Knowledge
  - Changing Data
  - Alternatives ~~which~~ That must be Evaluated
  - Procedural Complexity

©1985 by INPUT. Reproduction Prohibited.

EXHIBIT IV-6

EXPERT SYSTEM ~~SOFTWARE TOOLS~~

VENDOR	PRODUCT	HARDWARE	PRICE
Expert Systems, Inc.	Expert Ease	IBM PC, DEC Rainbow, Victor Technologies 9000	$125 $2,000
Teknowledge	M.1 S.1	IBM PC Xerox 1100 & 1108 and DEC VAX	$12,500 $50,000 - $80,000
Intellicorp	Knowledge Engineering Environment (KEE)	Symbolics 3600 Xerox 1100 Xerox 1108	$60,000
Texas Instruments	Personal Consultant	TI Professional Computer	$3,000
Inference Corp.	Advanced Reasoning Tool (ART)	Symbolics 3600 Lisp Machine DEC VAX	$48,000 - $60,000
General Research	The Intelligent Machine Model (TIMM)	DEC VAX IBM PC	$39,500 $9500

See M 84A Exhibit IV-3

©1985 by INPUT. Reproduction Prohibited.

EXHIBIT IV-7

CATEGORIES OF EXPERT SYSTEMS AND REPRESENTATIVE IMPACTS

CATEGORY	SOME POSSIBLE IMPACTED PERSONNEL
Interpretive Systems	Data Entry Personnel Business Analysts
Prediction Systems	Forecasters (Including Consultants)
Diagnosis Systems	Doctors Field Engineers
Design Systems	Systems and VLSI Designers Accounting and Financial Systems Personnel
Planning Systems	Programmers Schedulers
Monitoring Systems	Project and Financial Control Personnel
Debugging Systems	Programmers Circuit Designers
Repair Systems	Field Engineers
Instruction Systems	Teachers Training Personnel
Control Systems	Industrial Engineers Investment Analy~~sts~~ ysts

©1985 by INPUT. Reproduction Prohibited.